arXiv:1605.03636v1 [cs.PL] 11 May 2016

Tighter Loop Bound Analysis

(Technical report)

Pavel Čadek
Faculty of Informatics
Vienna University of Technology
Austria

Jan Strejček
Faculty of Informatics
Masaryk University, Brno
Czech Republic

Marek Trtík
LaBRI
University of Bordeaux
France

Abstract

We present a new algorithm for computing upper bounds on the number of executions of each program instruction during any single program run. The upper bounds are expressed as functions of program input values. The algorithm is primarily designed to produce bounds that are relatively tight, i.e. not unnecessarily blown up. The upper bounds for instructions allow us to infer loop bounds, i.e. upper bounds on the number of loop iterations. Experimental results show that the algorithm implemented in a prototype tool LOOPERMAN often produces tighter bounds than current tools for loop bound analysis.

Keywords: Loop bound, WCET, symbolic execution, path counter, loop summarisation, static analysis.

1 Introduction

The goal of *loop bound analysis* is to derive for each loop in a given program an upper bound on the number of its iterations during any execution of the program. These bounds can be parametrized by the program input. The loop bound analysis is an active research area with two prominent applications: *program complexity* analysis and *worst case execution time* (WCET) analysis.

The aim of program complexity analysis is to derive an asymptotic complexity of a given program. The complexity is commonly considered by programmers in their everyday work and it is also used in specifications of programming languages, e.g. every implementation of the standard template library of C++ has to have the prescribed complexities. Loop bound analysis clearly plays a crucial role in program complexity analysis. In this context, emphasis is put on large coverage of the loop bound analysis (i.e. it should find some bounds for as many program loops as possible), while there are only limited requirements on tightness of the bounds as asymptotic complexity is studied.

A typical application scenario for WCET analysis is to check whether a given part of some critical system finishes its execution within an allocated time budget. One step of the decision process is to compute loop bounds. Tightness of the bounds is very important here as an untight bound can lead to a spuriously negative answer of the analysis (i.e. 'the allocated time budged can be exceeded'), which may imply unnecessary additional costs, e.g. for system redesign or for hardware components with higher performance. The WCET analysis can also be used by schedulers to estimate the run-time of individual tasks.

The problem to infer loop bounds has recently been refined into the *reachability-bound problem* [16], where the goal is to find an upper bound on the number of executions of a given program instruction during any single run of a given program. One typically asks for a reachability bound on some resource demanding instruction like memory allocation. Reachability bound analysis is more challenging than loop bound analysis as, in order to get a reasonably precise bound, branching inside loops must be taken into account.

This paper presents a new algorithm that infers reachability bounds. More precisely, for each instruction of a given program, the algorithm tries to find an upper bound on the number of executions of the instruction in any single run of the program. The bounds are parametrized by the program input. The reachability bounds can be directly used to infer loop bounds and asymptotic program complexity. Our algorithm builds on *symbolic execution* [19] and *loop summarisation* adopted from [25]. In comparison with other techniques for reachability bound or loop bound analysis, our algorithm brings the following features:

- It utilizes a loop summarisation technique that computes precise values of program variables as functions of loop iteration counts.
- It distinguishes different branches inside loops and computes bounds for each of them separately.
- If more different bounds arise, it handles all of them while other techniques usually choose nondeterministically one of them.
- It can detect logarithmic bounds.
- Upper bounds for nested loops are computed more precisely: while other techniques typically multiply a bound for the outer loop by a maximal bound on iterations of the inner loop during one iteration of the outer loop, we sum the bounds for the inner loop over all iterations of the outer loop.

All these features have a positive effect on tightness of produced bounds.

We can explain the basic idea of our algorithm on the flowgraph on the right. The node a is the entry location, d is the exit location, and locations b, c form a loop. An initial value of `x` represents program input. We symbolically execute each path in the loop and assign an iteration counter to it. Then we try to express the effect of arbitrarily many iterations of the loop using the iteration counters as parameters. The loop in our example has just one path bcb that increments `i` by 2. Hence, the value of `i` after κ iterations is $\underline{i}' + 2\kappa$, where $\underline{i}'$ denotes the value of `i` before the loop execution starts. We combine this

```
void f(int x) {
   int i=5;
   while (i<x)
      i=i+2;
}
```

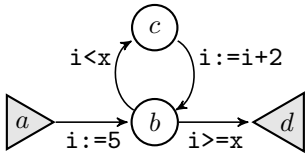

Figure 1: A simple C program with its flowgraph.

loop summary with the program state just before entering the loop, which is obtained by symbolic execution of the corresponding part of the program. In our example, we get that the value of `i` after κ iterations of the loop is $5+2\kappa$. To enter another iteration, the condition `i<x` must hold. If we replace the variables `i` and `x` by their current values, we get the condition $5 + 2\kappa < \underline{x}$, where $\underline{x}$ refers to the initial value of `x`. This condition is satisfied only if $\kappa < \frac{\underline{x}-5}{2}$. As κ is an iteration counter, it has to be a non-negative integer. Hence, we get the bound on the number of loop iterations $\mathbf{max}\{0, \lceil\frac{\underline{x}-5}{2}\rceil\}$, which is assigned to all edges in the loop. Edges outside the loop are visited at most once. The situation is more complicated if we have loops with more loop paths, nested loops, or loops where a run can cycle forever. The algorithm is described in Section 3.

We have implemented our algorithm in an experimental tool LOOPERMAN. Comparison with several leading loop bound analysis tools shows that our approach often provides tighter loop bounds. For example, our tool is currently the only one that detects that the inner cycle of the BubbleSort algorithm makes $\frac{n\cdot(n-1)}{2}$ iterations in total (i.e. during all iterations of the outer loop) when sorting an array of n elements, while other tools provide only the bound n^2 or $\mathcal{O}(n^2)$. Section 5 presents the comparison of LOOPERMAN with four other loop bound analysis tools on 199 benchmarks.

2 Preliminaries

In this section we introduce or recall some terminology intensively used in the following sections. For simplicity, this paper focuses on programs manipulating only integer scalar variables `a`, `b`, ... and read-only multidimensional integer array variables `A`, `B`, ..., and with no function calls.

2.1 Flowgraph, backbone, loop, induced flowgraph

An analysed program is represented as a *flowgraph* $P = (V, E, l_{beg}, l_{end}, \iota)$, where (V, E) is a finite oriented graph, $l_{beg}, l_{end} \in V$ are different *begin* and *end nodes* respectively. The in-degree of l_{beg} and the out-degree of l_{end} are 0, and $\iota : E \to \mathcal{I}$ labels each edge e by an *instruction* $\iota(e)$. The out-degree of all nodes except the end node is 1 or 2. Nodes with out-degree 2 are called *branching* nodes. We use two kinds of instructions: an assignment instruction `a:=`*expr* for some scalar variable `a` and some program expression *expr* over program variables, and an assumption `assume`(γ) for some quantifier-free formula γ over program variables. Out-edges of any branching node are labelled with assumptions `assume`(γ) and `assume`$(\neg\gamma)$ for some γ. We often omit the keyword `assume` in flowgraphs.

A *path* in a flowgraph is a (finite or infinite) sequence $\pi = v_1 v_2 \cdots$ of nodes such that $(v_i, v_{i+1}) \in E$ for all v_i, v_{i+1} in the sequence. Paths are denoted by Greek letters.

A *backbone* in a flowgraph is an acyclic path leading from the begin node to the end node.

Let π be a backbone with a prefix αv. There is a *loop* C with a *loop entry* v along π, if there exists a path $v\beta v$ such that no node of β appears in α. The loop C is then the smallest set containing all nodes of all such paths $v\beta v$.

Each *run* of the program corresponds to a path in the flowgraph starting at l_{beg} and such that it is either infinite, or it is finite and ends in l_{end}.[1] Every run follows some backbone: it can escape from the backbone in order to perform one or more iterations in a loop along the backbone, but once the last iteration in the loop is finished, the execution continues along the backbone again. We thus talk about a *run along* a backbone. We can compute the backbone for a finite run corresponding to a path ρ by the following procedure: If ρ is acyclic, then the backbone is directly ρ. Otherwise, we find the leftmost repeating node in ρ, remove the part of ρ between the first and the last occurrence of this node (including the last occurrence), and repeat the procedure. In other words, the backbone arises from ρ by removing all cycles in it. For an infinite run, we extend the procedure to remove the suffix starting just after the leftmost node that repeats infinitely often on the run. The procedure returns an acyclic path that is a prefix of some backbones. We can associate the run to any backbone with this prefix.

For a loop C with a loop entry v along a backbone π, a *flowgraph induced by the loop*, denoted as $P(C, v)$, is derived from the subgraph of the original flowgraph induced by C, where v is marked as the begin node, a fresh end node v' is added, and every transition $(u, v) \in E$ leading to v is redirected to v' (we identify the edge (u, v') with (u, v) in the context of the original program). Each single iteration (i.e. a path from v back to v without visiting v inside) of the loop corresponds to a run of the induced flowgraph.

Note that program representation by flowgraphs and the above definition of loops easily handle many features of programming languages like `break`, `continue`, or even program loops constructed using `goto` instructions.

2.2 Symbolic execution

The basic idea of *symbolic execution* [19] is to replace input data of a program by *symbols* representing arbitrary data. Executed instructions then manipulate symbolic expressions over the symbols instead of exact values. A *symbolic expression* is any term of the theory of integers extended with functions **max** and **min**, rounding functions $\lceil\cdot\rceil$ and $\lfloor\cdot\rfloor$ for constant expressions over reals, and

- for each scalar variable `a`, an uninterpreted constant $\underline{a}$, which is a *symbol* representing any initial (input) value of the variable `a`,
- for each array variable `A`, an uninterpreted function $\underline{A}$ of the same arity as `A`, which is a *symbol* representing any initial (input) content of the array `A`,

[1] We assume that crashes or other undefined behaviour of program expressions are prevented by safety guards. For example, an expression `a/b` is guarded by `assume(b` $\neq 0$).

- a countable set $\{\kappa_1, \kappa_2, \ldots\}$ of artificial variables (not appearing in analysed programs), called *path counters* and ranging over non-negative integers,
- a special symbol $\star$ called *unknown*, and
- for each formula ψ build over symbolic expressions and two symbolic expressions ϕ_1, ϕ_2, a construct $\mathbf{ite}(\psi, \phi_1, \phi_2)$ of meaning "**i**f-**t**hen-**e**lse", that evaluates to ϕ_1 if ψ holds, and to ϕ_2 otherwise.

Let ψ, ϕ be symbolic expressions and x be a symbol or a path counter. By $\psi[x/\phi]$ we denote the expression ψ, where all occurrences of x are simultaneously replaced by ϕ. We further extend this notation such that $\psi[x_i/\phi_i \mid i \in I]$ denotes the expression ψ, where all occurrences of x_i are simultaneously replaced by ϕ_i, for each $i \in I$.

We sometimes add the upper index $\vec{\kappa}$ to expression names to denote that the expressions can contain path counters $\vec{\kappa} = (\kappa_1, \ldots, \kappa_n)$. We say that an expression is *$\vec{\kappa}$-free* if it contains no path counters.

Symbolic execution stores variable values in *symbolic memory* and all executable program paths are uniquely identified by corresponding *path conditions*. Here we provide brief descriptions of these terms. For more information see [19].

A *symbolic memory* is a function θ assigning to each scalar variable `a` a symbolic expression and to each array variable `A` the symbol $\underline{A}$ (the value of every array variable is always identical to its initial value as we consider programs with read-only arrays). A symbolic memory θ is called *initial*, if $\theta(\mathtt{a}) = \underline{a}$, for each scalar or array variable `a`.

We overload the notation $\theta(\cdot)$ to program expressions as follows. Let *expr* be a program expression over program variables $\mathtt{a}_1, \ldots, \mathtt{a}_n$. Then $\theta(expr)$ represents a symbolic expression obtained from *expr* such that we simultaneously replace all occurrences of the variables $\mathtt{a}_1, \ldots, \mathtt{a}_n$ by symbolic expressions $\theta(\mathtt{a}_1), \ldots, \theta(\mathtt{a}_n)$ respectively.

Symbolic execution of a path in a flowgraph starts with the initial symbolic memory and the memory is updated on assignments. For example, if the first executed assignment is `a:=2*a+b`, the initial symbolic memory θ is updated to the symbolic memory θ' where $\theta'(\mathtt{a}) = \theta(\mathtt{2*a+b}) = 2\underline{a} + \underline{b}$. If we later update θ' on `a:=1-a`, we get the memory θ'' such that $\theta''(\mathtt{a}) = \theta'(\mathtt{1-a}) = 1 - (2\underline{a} + \underline{b}) = 1 - 2\underline{a} - \underline{b}$.

If ψ is a symbolic expression over symbols $\{\underline{a}_i \mid i \in I\}$ corresponding to program variables $\{\mathtt{a}_i \mid i \in I\}$ respectively, then $\theta\langle\psi\rangle$ denotes the symbolic expression $\psi[\underline{a}_i/\theta(\mathtt{a}_i) \mid i \in I]$. For example, if $\theta(\mathtt{a}) = \kappa_1$ and $\theta(\mathtt{b}) = \underline{a} - \kappa_2$, then $\theta\langle 2\underline{a} + \underline{b}\rangle = 2\theta(\mathtt{a}) + \theta(\mathtt{b}) = 2\kappa_1 + \underline{a} - \kappa_2$. We apply the notation $\theta\langle\varphi\rangle$ and $\varphi[x/\psi]$ with the analogous meanings also to formulae φ built over symbolic expressions.

Note that $\theta_1\langle\theta_2(\mathtt{a})\rangle$ returns the value of `a` after a code with effect θ_1 followed by a code with effect θ_2. For example, if $\theta_1(\mathtt{a}) = 2\underline{a} + 1$ represents the effect of assignment `a:=2*a+1` and $\theta_2(\mathtt{a}) = \underline{a} - 2$ the effect of assignment `a:=a-2`, then $\theta_1\langle\theta_2(\mathtt{a})\rangle = \theta_1\langle\underline{a} - 2\rangle = (2\underline{a} + 1) - 2$ represents the effect of the two assignments applied in the sequence.

Given a path in a flowgraph leading from its begin node l_{beg}, the *path condition* is a formula over symbols, which determines the set of all possible program inputs for which the program execution will follow the path. A path condition is constructed during symbolic execution of the path. Initially, the path condition

is set to *true* and it can only be updated when an `assume`(γ) is executed. For example, if a symbolic execution reaches `assume(a > 5)` with a path condition φ and a symbolic memory $\theta(\mathtt{a}) = 2\underline{a} - 1$, then it updates the path condition to $\varphi \wedge 2\underline{a} - 1 > 5$.

2.3 Upper bound

An *upper bound* for an edge e in a flowgraph P is a $\vec{\kappa}$-free symbolic expression ρ such that whenever P is executed on any input, the instruction on edge e is executed at most ρ' times, where ρ' is the expression that we get by replacing each variable symbol by the input value of the corresponding variable.

3 The Algorithm

Recall that every program run follows some backbone and the run can diverge from the backbone only to loops along the backbone. The algorithm first detects all backbones. For each backbone π_i and each edge e, it computes a set of upper bounds $\beta_i(e)$ on the number of visits of the edge by any run following the considered backbone. As all these bounds are valid, the set $\beta_i(e)$ can be interpreted as a single bound $\mathbf{min}\,\beta_i(e)$ on visits of edge e by any run along π_i. At the end, the overall upper bound for an edge e can be computed as the maximum of these bounds over all backbones, i.e. as $\mathbf{max}\{\mathbf{min}\,\beta_i(e) \mid \pi_i \text{ is a backbone}\}$.

The algorithm consists of the following four procedures:

`executeProgram` is the starting procedure of the whole algorithm. It gets a flowgraph and computes all its backbones. Then it symbolically executes each backbone and computes for each edge a set of upper bounds on the number of visits of the edge by a run following the backbone. Whenever the symbolic execution visits a loop entry node, the procedure `processLoop` is called to get upper bounds on visits during loop execution.

`processLoop` gets a loop represented by the program induced by the loop. Note that each run of the induced program corresponds to one iteration of the loop and it follows some backbone of the induced program (the backbones are called *loop paths* in this context). The procedure then symbolically executes each loop path by recursive call of `executeProgram` (the nesting of recursive calls thus directly corresponds to the nesting of loops in the program). The recursive call of `executeProgram` produces, for each loop path, a symbolic memory and a path condition capturing the effect of a single iteration along the loop path. The procedure `processLoop` then calls `computeSummary`, which takes the symbolic memories after single loop iterations, assigns to each loop path a unique path counter κ_i, and computes a *parametrized symbolic memory* $\theta^{\vec{\kappa}}$ describing the effect of an arbitrary number of loop iterations. This symbolic memory is parametrized by path counters $\vec{\kappa} = (\kappa_1, \ldots, \kappa_k)$ representing the numbers of iterations along the corresponding loop paths. From the parametrized symbolic memory and from the path conditions corresponding to single loop iterations (received from the recursive call of `executeProgram`), we derive a *parametrized necessary condition* for each loop path, which is a formula over symbols and path counters $\vec{\kappa}$ that has to be satisfied when another loop iteration

along the corresponding loop path can be performed after $\vec{\kappa}$ loop iterations. Finally, `processLoop` infers upper bounds from these parametrized necessary conditions with the help of the procedure `computeBounds`.

`computeSummary` is a subroutine of `processLoop` that gets symbolic memories corresponding to one loop iteration along each loop path and it produces the parametrized symbolic memory $\theta^{\vec{\kappa}}$ after an arbitrary number of loop iterations (as mentioned above).

`computeBounds` is another subroutine of `processLoop`. It gets a set I of loop paths and the corresponding parametrized necessary conditions and tries to derive some upper bounds on the number of loop iterations along loop paths from I.

We describe the four procedures in the following four subsections. The procedure `processLoop` is described as the last one as it calls the other three procedures.

3.1 Algorithm `executeProgram`

The procedure `executeProgram` of Algorithm 1 takes a flowgraph as input, determines its backbones, and symbolically executes each backbone separately. For a backbone π_i, symbolic execution computes symbolic memory θ_i, path condition φ_i, and *bound function* β_i assigning to each edge e a set of symbolic expressions that are valid upper bounds on the number of visits of edge e during any single run along the backbone. Each such a set $\beta_i(e)$ of bounds could be replaced by a single bound $\mathbf{min}\ \beta_i(e)$, but we prefer to keep it as a set of simpler expressions to increase the success rate of expression matching in procedure `processLoop` (we point out the reason in Section 3.4).

The symbolic execution proceeds in the standard way until we enter a loop entry (line 9). Then we call procedure `processLoop` on the loop, current symbolic memory and path condition. The procedure returns function β_{loop} of upper bounds on visits of loop edges during execution of the loop, and a symbolic memory after execution of the loop. We add these bounds and the former bounds in the **foreach** loop at line 12 and continue the execution along the backbone. If the `processLoop` procedure cannot determine the value of some variable after the loop, it simply uses the symbol $\star$ (unknown).

Another difference from the standard symbolic execution is at line 14 where we suppress insertion of predicates containing $\star$ to the path condition. As a consequence, a path condition of our approach is no longer a necessary and sufficient condition on input values to lead the program execution along the corresponding path (which is the case in standard symbolic execution), but it is only a necessary condition on input values of a run to follow the backbone.

After processing an edge of the backbone, we increase the corresponding bounds by one (line 18).

At the end of the procedure, the resulting bounds for each edge are computed as the maximum of previously computed bounds for the edge over all backbones (see the **foreach** loop at line 20). Besides these bounds, the procedure also returns each backbone with the symbolic memory and path condition after its execution.

Algorithm 1: `executeProgram`

Input:
P *// a flowgraph*
Output:
$\{(\pi_1, \theta_1, \varphi_1), \ldots, (\pi_k, \theta_k, \varphi_k)\}$ *// backbones π_i (with symbolic memory θ_i // and path condition φ_i after execution // along π_i)*
β *// for each edge e of P, $\beta(e)$ is a set of upper bounds for e*

1 `states` $\longleftarrow \emptyset$
2 Compute the set of backbones $\{\pi_1, \ldots, \pi_k\}$ of P.
3 **foreach** $i = 1, \ldots, k$ **do**
4 Initialize θ_i to return $\underline{a}$ for each (scalar or array) variable `a`.
5 $\varphi_i \longleftarrow$ *true*
6 Initialize β_i to return $\{0\}$ for each edge.
7 Let $\pi_i = v_1 \ldots v_n$.
8 **foreach** $j = 1, \ldots, n-1$ **do**
9 **if** v_j *is a loop entry* **then**
10 Let C be the loop with the loop entry v_j along π_i.
11 $(\beta_{loop}, \theta_i) \longleftarrow$ `processLoop`$(P(C, v_j), \theta_i, \varphi_i)$
12 **foreach** *edge e of $P(C, v_j)$* **do**
13 $\beta_i(e) \longleftarrow \{\rho_1 + \rho_2 \mid \rho_1 \in \beta_i(e),\ \rho_2 \in \beta_{loop}(e)\}$
14 **if** $\iota((v_j, v_{j+1}))$ *has the form* `assume`(ψ) *and $\theta_i(\psi)$ contains no $\star$* **then**
15 $\varphi_i \longleftarrow \varphi_i \wedge \theta_i(\psi)$
16 **if** $\iota((v_j, v_{j+1}))$ *has the form* `a` $:=$ *expr* **then**
17 $\theta_i($`a`$) \longleftarrow \theta_i(expr)$
18 $\beta_i((v_j, v_{j+1})) \longleftarrow \{\rho + 1 \mid \rho \in \beta_i((v_j, v_{j+1}))\}$
19 Insert $(\pi_i, \theta_i, \varphi_i)$ into `states`.
20 **foreach** *edge e of P* **do**
21 $\beta(e) \longleftarrow \{\mathbf{max}\{\rho_1, \ldots, \rho_k\} \mid \rho_1 \in \beta_1(e), \ldots,\ \rho_k \in \beta_k(e)\}$
22 **return** (`states`, β)

3.2 Algorithm `computeSummary`

The procedure `computeSummary` of Algorithm 2 is a subpart of the procedure of the same name introduced in [25].

The procedure gets loop paths $\pi_1, \ldots, \pi_l$ together with symbolic memories $\theta_1, \ldots, \theta_l$, where each θ_i represents the effect of a single iteration along π_i. Then it assigns fresh path counters $\vec{\kappa} = (\kappa_1, \ldots, \kappa_l)$ to the loop paths and computes the parametrized symbolic memory $\theta^{\vec{\kappa}}$ after $\vec{\kappa}$ iterations of the loop, i.e. after $\sum_{1 \le i \le l} \kappa_i$ iterations where exactly κ_i iterations follow π_i for each i and there is no assumption on the order of iterations along different loop paths.

Note that the value of some variables can depend on the order of iterations along different loop paths. If we do not find the precise value of some variable after $\vec{\kappa}$ iterations, then $\theta^{\vec{\kappa}}$ assigns $\star$ (unknown) to this variable.

To be on safe side, we start with $\theta^{\vec{\kappa}}$ assigning $\star$ to all scalar variables. Then we gradually improve the precision of $\theta^{\vec{\kappa}}$ as long as there is some progress. The crucial step is the computation of an improved value b for a scalar variable `a` at line 6. The value b is defined as $\star$ except the following four cases.

1. For each loop path π_i, we have $\theta_i($`a`$) = \underline{a}$. In other words, the value of `a` is not changed in any iteration of the loop. This case is trivial. We set

Algorithm 2: `computeSummary`

Input:
$\{(\pi_1, \theta_1), \ldots, (\pi_l, \theta_l)\}$ *// each θ_i is a memory after a single execution of π_i*
Output:
$\theta^{\vec{\kappa}}$ *// the symbolic memory after $\vec{\kappa}$ iterations of backbones $\pi_1, \ldots, \pi_l$*

```
1  Introduce fresh path counters κ⃗ = (κ_1, ..., κ_l) for backbones π_1, ..., π_l resp.
2  Initialize θ^κ⃗ to return ⋆ for each scalar variable, and A for each array variable A.
3  repeat
4      change ⟵ false
5      foreach variable a such that θ^κ⃗(a) = ⋆ do
6          Compute an improved value b for the variable a from symbolic
           memories θ_1, ..., θ_l and θ^κ⃗.
7          if b ≠ ⋆ then
8              θ^κ⃗(a) ⟵ b
9              change ⟵ true
10 until change = false
11 return θ^κ⃗
```

$b = \underline{a}$.

2. For each loop path π_i, either $\theta_i(\mathtt{a}) = \underline{a}$ or $\theta_i(\mathtt{a}) = \underline{a} + d_i$ (resp. $\theta_i(\mathtt{a}) = \underline{a} \cdot d_i$) for some symbolic expression d_i such that $\theta^{\vec{\kappa}}\langle d_i \rangle$ contains neither $\star$ nor any path counters. Let us assume that the latter possibility holds for loop paths $\pi_1, \ldots, \pi_m$. The condition on $\theta^{\vec{\kappa}}\langle d_i \rangle$ guarantees that the value of d_i is constant during all iterations over the loop. In this case, we set $b = \underline{a} + \sum_{i=1}^{m} \theta^{\vec{\kappa}}\langle d_i \rangle \cdot \kappa_i$ (resp. $b = \underline{a} \cdot \prod_{i=1}^{m} \theta^{\vec{\kappa}}\langle d_i \rangle^{\kappa_i}$).

3. There exists a symbolic expression d such that $\theta^{\vec{\kappa}}\langle d \rangle$ contains neither $\star$ nor any path counters. For each loop path π_i, either $\theta_i(\mathtt{a}) = \underline{a}$ or $\theta_i(\mathtt{a}) = d$. Let us assume that the latter possibility holds for loop paths $\pi_1, \ldots, \pi_m$. In other words, the value of $\mathtt{a}$ is set to d in each iteration along loop path π_i for $1 \leq i \leq m$, while it is unchanged in any other iteration. Hence, we set $b = \mathbf{ite}(\sum_{1 \leq i \leq m} \kappa_i > 0, \theta^{\vec{\kappa}}\langle d \rangle, \underline{a})$.

4. For one loop path, say π_i, $\theta_i(\mathtt{a}) = d$ for some symbolic expression d such that $\theta^{\vec{\kappa}}\langle d \rangle$ contains neither $\star$ nor any path counters except κ_i. Further, for each loop path π_j such that $i \neq j$, $\theta_j(\mathtt{a}) = \underline{a}$. That is, only iterations along π_i modify $\mathtt{a}$ and they set it to a value independent on other path counters than κ_i. Note that if we assign d to $\mathtt{a}$ in the κ_i-th iteration along π_i, then the actual assigned value of d is the value after $\kappa_i - 1$ iterations along the paths. Therefore we set $b = \mathbf{ite}(\kappa_i > 0, (\theta^{\vec{\kappa}}\langle d \rangle)[\kappa_i / \kappa_i - 1], \underline{a})$.

Note that one can add another cases covering other situations where the value of $\mathtt{a}$ can be expressed precisely, e.g. the case capturing geometric progressions. Also note that the value of any variable in the resulting symbolic memory $\theta^{\vec{\kappa}}$ is expressed either precisely, or it is unknown (i.e. $\star$).

Algorithm 3: `computeBounds`

Input:
I // *indices of backbones*
φ // *a necessary condition to perform an iteration along a backbone with an index in I after $\vec{\kappa}$ iterations*

Output:
B // *upper bounds on the number of iterations along backbones with indices in I*

1 **if** $\varphi[\kappa_i/0 \mid i \in I]$ *is not satisfiable* **then return** $\{0\}$ $B \longleftarrow \emptyset$
2 **foreach** inequality $\sum_{j \in J \supseteq I} a_j \kappa_j < b$ implied by φ, where each a_j is a positive integer and b is $\vec{\kappa}$-free **do**
3 $\quad B \longleftarrow B \cup \{\mathbf{max}\{0, \lceil b / \mathbf{min}\{a_i \mid i \in I\}\rceil\}\}$
4 **return** B

3.3 Algorithm `computeBounds`

The procedure `computeBounds` of Algorithm 3 gets a set I of selected loop path indices, and a necessary condition φ to perform an iteration along some loop path with an index in I (we talk about an *iteration along* I for short) after $\vec{\kappa}$ previous loop iterations. From this information, the procedure infers upper bounds on the number of loop iterations along I.

We would like to find a tight upper bound, i.e. a $\vec{\kappa}$-free symbolic expression B such that there exist some values of symbols (given by a valuation function v) for which the necessary condition $\varphi[\,\underline{a}/v(\underline{a}) \mid \underline{a}$ is a symbol$\,]$ to make another iteration along I is satisfiable whenever the number of finished iterations along I is less than $B[\,\underline{a}/v(\underline{a}) \mid \underline{a}$ is a symbol$\,]$ and the same does not hold for the expression $B + 1$. An effective algorithm computing these tight bounds is an interesting research topic itself.

The presented procedure infers some bounds only for two special cases. Line 1 covers the case when even the first iteration along any loop path in I is not possible: the procedure then returns the bound 0.

The other special case is the situation when the necessary condition implies an inequality of the form $\sum_{j \in J \supseteq I} a_j \kappa_j < b$, where each a_j is a positive integer and b is $\vec{\kappa}$-free. To detect these cases, we transform the necessary condition to the conjunctive normal form and look for clauses that contain just one predicate and try to transfer the predicate into this form by basic arithmetical operations knowing that φ holds. Each such inequality implies the following:

$$\sum_{j \in J \supseteq I} a_j \kappa_j < b \implies \sum_{i \in I} a_i \kappa_i < b \implies \mathbf{min}\{a_i \mid i \in I\} \cdot \sum_{i \in I} \kappa_i < b.$$

Hence, $\sum_{i \in I} \kappa_i < \lceil b / \mathbf{min}\{a_i \mid i \in I\}\rceil$ has to be satisfied to perform another iteration along I after $\vec{\kappa}$ previous iterations including $\sum_{i \in I} \kappa_i$ iterations along I. As all path counters are non-negative integers, we derive the bound $\mathbf{max}\{0, \lceil b / \mathbf{min}\{a_i \mid i \in I\}\rceil\}$ on iterations along I.

3.4 Algorithm `processLoop`

The procedure `processLoop` of Algorithm 4 gets a flowgraph Q representing the body of a loop, i.e. each run of Q corresponds to one iteration of the original

Algorithm 4: `processLoop`

Input:
Q // *a flowgraph induced by a loop*
θ_{in} // *a symbolic memory when entering the loop*
φ_{in} // *a path condition when entering the loop*

Output:
β_{loop} // *upper bounds for all edges in the loop*
θ_{out} // *symbolic memory after the loop*

1 Initialize β_{loop} to return $\emptyset$ for each edge e of Q.
2 $(\{(\pi_1, \theta_1, \varphi_1), \ldots, (\pi_k, \theta_k, \varphi_k)\}, \beta_{inner}) \longleftarrow$ `executeProgram(`Q`)`
3 $\theta^{\vec{\kappa}} \longleftarrow$ `computeSummary(`$\{(\pi_1, \theta_1), \ldots, (\pi_k, \theta_k)\}$`)`
4 $\varphi_i^{\vec{\kappa}} \longleftarrow \varphi_{in} \wedge \theta_{in}\langle\theta^{\vec{\kappa}}\langle\varphi_i\rangle\rangle$ for each $i \in \{1, \ldots, k\}$
5 $\beta^{\vec{\kappa}}(e) \longleftarrow \{\theta_{in}\langle\theta^{\vec{\kappa}}\langle\rho\rangle\rangle \mid \rho \in \beta_{inner}(e)\}$ for each edge e of Q
6 **foreach** edge e of Q **do**
7 $I \longleftarrow \{i \mid e$ is on π_i or on a loop along $\pi_i\}$
8 $B_{outer} \longleftarrow$ `computeBounds(`$I, \bigvee_{i \in I} \varphi_i^{\vec{\kappa}}$`)`
9 **if** $0 \in B_{outer}$ **then**
10 $\beta_{loop}(e) \longleftarrow \{0\}$
11 **else**
12 **foreach** $\rho_{inner} \in \beta^{\vec{\kappa}}(e)$ **do**
13 **if** $\rho_{inner} \equiv c$ where c is $\vec{\kappa}$-free **then**
14 $\beta_{loop}(e) \longleftarrow \beta_{loop}(e) \cup \{c \cdot \rho_{outer} \mid \rho_{outer} \in B_{outer}\}$
15 **else if** $\rho_{inner} \equiv \mathbf{max}\{c, b + \sum_{i=1}^{k} a_i\kappa_i\}$ where c, b and all a_i are $\vec{\kappa}$-free **then**
16 $J \longleftarrow \{j \mid j \notin I \wedge a_j \neq 0\}$
17 $B_J \longleftarrow$ `computeBounds(`$J, \bigvee_{j \in J} \varphi_j^{\vec{\kappa}}$`)`
18 **if** $B_J \neq \emptyset$ **then**
19 $b' \longleftarrow b + \mathbf{max}\{0, a_j \mid j \in J\} \cdot \mathbf{min}\, B_J$
20 $a \longleftarrow \mathbf{max}\{a_i \mid i \in I\}$
21 **foreach** $\rho_{outer} \in B_{outer}$ **do**
22 $\beta_{loop}(e) \longleftarrow \beta_{loop}(e) \cup \{\sum_{K=0}^{\rho_{outer}-1} \mathbf{max}\{c, b' + a \cdot K\}\}$
23 **foreach** edge e of Q **do**
24 **if** an edge e' of Q such that $\beta^{\vec{\kappa}}(e') = \emptyset$ is reachable from e in Q **then**
25 $\beta_{loop}(e) \longleftarrow \{\rho_1 + \rho_2 \mid \rho_1 \in \beta_{loop}(e),\ \rho_2 \in \beta^{\vec{\kappa}}(e)$, and ρ_2 is $\vec{\kappa}$-free$\}$
26 $\theta_{out}(\mathtt{a}) \longleftarrow \theta_{in}\langle\theta^{\vec{\kappa}}(\mathtt{a})\rangle$ for each variable `a`
27 Eliminate $\vec{\kappa}$ from θ_{out}.
28 **return** $(\beta_{loop}, \theta_{out})$

loop. We symbolically execute Q using the recursive call of `executeProgram` at line 2. We obtain all loop paths $\pi_1, \ldots, \pi_k$ of Q and bounds β_{inner} on visits of each edge in the loop during any single iteration of the loop. For each π_i, we also get symbolic memory θ_i after one iteration along π_i and a necessary condition φ_i to perform this iteration. The procedure `computeSummary` produces the parametrized symbolic memory $\theta^{\vec{\kappa}}$ after $\vec{\kappa}$ iterations. Symbols $\underline{a}$ appearing in $\theta^{\vec{\kappa}}$ refer to variable values before the loop is entered. If we combine $\theta^{\vec{\kappa}}$ with the symbolic memory before entering the loop θ_{in}, we get the symbolic memory after execution of the code preceding the loop and $\vec{\kappa}$ iterations of the loop. We use this combination to derive necessary conditions $\varphi_i^{\vec{\kappa}}$ to perform another iteration along π_i and upper bounds $\beta^{\vec{\kappa}}$ on visits of loop edges in the next iteration of

the loop.

The **foreach** loop at line 6 computes upper bounds for all edges of the processed loop on visits during all its complete iterations (incomplete iterations when a run cycles in some nested loop forever are handled later). We already have the bounds $\beta^{\vec{\kappa}}$ on visits in a single iteration after $\vec{\kappa}$ preceding iterations. For each edge e, we compute the set I of all loop path indices such that iterations along these loop paths can visit e. The `computeBounds` procedure at line 8 takes I and a necessary condition to perform an iteration along I after $\vec{\kappa}$ iterations and computes bounds B_{outer} on the number of iterations along I. If there is 0 among these bounds, e cannot be visited by any complete iteration and the computation for e is over. Otherwise we try to compute some overall bounds for each bound ρ_{inner} on the visits of e during one iteration (after $\vec{\kappa}$ iterations) separately. If ρ_{inner} is a $\vec{\kappa}$-free expression (line 13), then it is constant in each iteration and we simply multiply it with every bound on the number of iterations along I. The situation is more difficult if ρ_{inner} contains some path counters. We can handle the frequent case when it has the form $\mathbf{max}\{c, b + \sum_{i=1}^{k} a_i\kappa_i\}$, where $a_1, \ldots, a_k, b, c$ are $\vec{\kappa}$-free (see line 15 and note that this is the reason for keeping the bounds simple). First we get rid of path counters κ_j that have some influence on this bound (i.e. $a_j \neq 0$), but e cannot be visited by any iteration along loop path π_j. Let J be the set of indices of such path counters (line 16). We try to compute bounds B_J on the number of iterations along J (line 17), which are also the bounds on $\sum_{j\in J} \kappa_j$. Note that if $J = \emptyset$, we call `computeBounds(`$\emptyset, \mathit{false}$`)`, which immediately returns $\{0\}$. If we get some bounds in B_J, we can overapproximate $\sum_{i=1}^{k} a_i\kappa_i$ as follows:

$$\sum_{i=1}^{k} a_i\kappa_i = \sum_{j\in J} a_j\kappa_j + \sum_{i\in I} a_i\kappa_i \leq$$
$$\leq \ \mathbf{max}\{0, a_j \mid j \in J\} \cdot \mathbf{min}\, B_J + \mathbf{max}\{a_i \mid i \in I\} \cdot \sum_{i\in I} \kappa_i$$

Using the definitions of b' and a at lines 19–20, we overapproximate the bound ρ_{inner} on visits of e during one iteration along I after $\vec{\kappa}$ loop iterations by

$$\rho_{inner} = \mathbf{max}\{c, b + \sum_{i=1}^{k} a_i\kappa_i\} \ \leq \ \mathbf{max}\{c, b' + a \cdot \sum_{i\in I} \kappa_i\}.$$

As K-th iteration along I is preceded by $K - 1$ iterations along I, the edge e can be visited at most $\mathbf{max}\{c, b' + a \cdot (K - 1)\}$ times during K-th iteration. For each bound ρ_{outer} on the iterations along I, we can now compute the total bound on visits of e as $\sum_{K=0}^{\rho_{outer}-1} \mathbf{max}\{c, b' + a \cdot K\}$.

Until now we have considered visits of loop edges during *complete* iterations. However, it may also happen that an iteration is started, but never finished because the execution keeps looping forever in some nested loop. For example, in the program `while(x>0){x:=x-1;while(true){}}`, we easily compute bound 0 on the number of complete iterations of the outer loop and thus we assign bound 0 to all loop edges at line 10. However, some edges of the loop are visited. These incomplete iterations are treated by the **foreach** loop at line 23. Whenever an edge e can be visited by an incomplete iteration (which is detected by existence of some subsequent edge e' without any bound and thus potentially lying on

```
int nonzeros(int n, int* A) {
  int k=0;
  for (int i=0; i<n && k<3; i++)
    if (A[i]!=0)
      k++;
  return k;
}
```

(a)

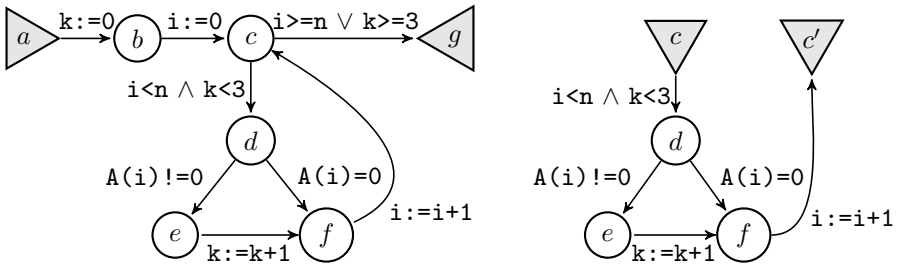

(b) (c)

Figure 2: C program (a) with the corresponding flowgraph (b), and its induced flowgraph $P(\{c, d, e, f\}, c)$ (c).

an infinite nested loop), we add the ($\vec{\kappa}$-free) bounds on visits of e during one iteration to the total bounds for e. If there is no such κ-free bound, we leave e unbounded to be on the safe side.

Finally, the lines 26 and 27 combine the symbolic memory before the loop with the effect of the loop and eliminate loop counters from the resulting symbolic memory θ_{out}. Roughly speaking, the elimination replaces every expression that is not $\vec{\kappa}$-free by $\star$. In fact, the elimination can be done in a smarter way. For example, after the loop in the program `i:=0;while(i<n){i:=i+1}`, the elimination can replace κ by $\mathbf{max}\{0, \theta_{out}(\mathtt{n})\}$.

4 Example

We execute the proposed technique on the example of Figure 2(a–b). The program accepts an array `A` of size `n` and it counts up to three non-zero elements in the array.

We will observe a run of our algorithm on the flowgraph step by step. We will follow calls to individual procedures where we present the current state of the computation in terms of variables of the procedure at the top of the call stack. In descriptions of symbolic memories we omit variables that are not modified in the program: the value of `n` and `A` is always $\underline{n}$ and $\underline{A}$, respectively.

The execution starts by calling `executeProgram` with the flowgraph at Figure 2(b). The flowgraph has only one backbone $\pi_1 = abcg$. The node c is the loop entry to the loop $\{c, d, e, f\}$ along the backbone. Symbolic execution of π_1 up to c is straightforward and leads to the symbolic memory $\theta_1(\mathtt{k}) = \theta_1(\mathtt{i}) = 0$, the path condition $\varphi_1 = \mathit{true}$, and the bound function β_1 maps each edge to $\{0\}$ except $\beta_1((a, b)) = \beta_1((b, c)) = \{1\}$. At the entry node c we build

an induced flowgraph $P(\{c,d,e,f\},c)$ depicted in Figure 2(c). Then we call `processLoop`$(P(\{c,d,e,f\},c),\theta_1,\varphi_1)$.

`processLoop` calls `executeProgram` with the flowgraph at Figure 2(c). Since the flowgraph does not contain any loop, `executeProgram` symbolically executes both loop paths $cdefc'$ and $cdfc'$ and returns the result to the caller. `processLoop` thus receives backbones $\pi_1 = cdefc'$ and $\pi_2 = cdfc'$, the corresponding symbolic memories and path conditions

$$\begin{aligned} \theta_1(\mathtt{i}) &= \underline{i}+1 & \theta_2(\mathtt{i}) &= \underline{i}+1 \\ \theta_1(\mathtt{k}) &= \underline{k}+1 & \theta_2(\mathtt{k}) &= \underline{k} \\ \varphi_1 &= \underline{i} < \underline{n} \ \wedge \ \underline{k} < 3 \ \wedge & \varphi_2 &= \underline{i} < \underline{n} \ \wedge \ \underline{k} < 3 \ \wedge \\ & \quad \underline{A}(\underline{i}) \neq 0 & & \quad \underline{A}(\underline{i}) = 0 \end{aligned}$$

and a bound function β_{inner} assigning $\{1\}$ to each edge of the flowgraph. Now we call `computeSummary` for the symbolic memories θ_1 and θ_2 and we get path counters $\vec{\kappa} = (\kappa_1,\kappa_2)$ for the backbones (π_1,π_2) respectively and the following symbolic memory $\theta^{\vec{\kappa}}$

$$\theta^{\vec{\kappa}}(\mathtt{i}) = \underline{i}+\kappa_1+\kappa_2 \qquad \theta^{\vec{\kappa}}(\mathtt{k}) = \underline{k}+\kappa_1$$

Next, at line 4 we compute necessary conditions (to perform another iteration along backbones π_1 and π_2 respectively)

$$\begin{aligned} \varphi_1^{\vec{\kappa}} &= \kappa_1+\kappa_2 < \underline{n} \ \wedge \ \kappa_1 < 3 \ \wedge \ \underline{A}(\kappa_1+\kappa_2) \neq 0 \\ \varphi_2^{\vec{\kappa}} &= \kappa_1+\kappa_2 < \underline{n} \ \wedge \ \kappa_1 < 3 \wedge \ \underline{A}(\kappa_1+\kappa_2) = 0 \end{aligned}$$

The next line produces bound function $\beta^{\vec{\kappa}}$ which is the same as β_{inner}, in this case. Now we have all data we need to start the computation of resulting bounds for all five edges of the passed flowgraph.

The main part of this computation is performed in the loop at line 6. We show the computation for the edge (e,f). First we call `computeBounds`$(\{1\},\varphi_1^{\vec{\kappa}})$. Since the passed formula is satisfiable for $\kappa_1 = 0$, we get to the line 2. The condition there is satisfied for predicates $\kappa_1+\kappa_2 < \underline{n}$ and $\kappa_1 < 3$. The computation on the subsequent line is straightforward for both predicates, so we receive the set $B_{outer} = \{\mathbf{max}\{0,\underline{n}\},\mathbf{max}\{0,3\}\} = \{\mathbf{max}\{0,\underline{n}\},3\}$. Then we get to the line 14 in `processLoop` (because $\beta^{\vec{\kappa}}((e,f)) = \{1\}$), where we receive $\beta_{loop}((e,f)) = \{\mathbf{max}\{0,\underline{n}\},3\}$. The computation proceeds similarly for other edges, but for (c,d), (d,f), (f,c) it produces only one bound $\{\mathbf{max}\{0,\underline{n}\}\}$. Since, the condition at line 24 is false for all edges, the resulting β_{loop} returns $\{\mathbf{max}\{0,\underline{n}\},3\}$ for (d,e) and (e,f), and $\{\mathbf{max}\{0,\underline{n}\}\}$ for the others. The resulting symbolic memory θ_{out} assigns $\star$ to `i` and `k`.

The control-flow then returns back to `executeProgram` where we update β_1 according to received β_{loop}. Then we symbolically execute the remaining edge (c,g). The computation in the loop at line 20 computes maximum over all bounds for a considered edge. The algorithm then terminates with the bound function β assigning $\{1\}$ to edges (a,b), (b,c), (c,g), the set $\{\mathbf{max}\{0,\underline{n}\}\}$ to edges (c,d), (d,f), (f,c), and the set $\{\mathbf{max}\{0,\underline{n}\},3\}$ to (d,e) and (e,f).

We can conclude for the program in Figure 2(a) that the loop can be executed only if the program is called with some positive integer $\underline{n}$ for the parameter `n`. In that case the loop is executed at most $\mathbf{max}\{0,\underline{n}\}$ times (according

```
void bubble_sort(int n, int* A) {
  for(int i = 0; i < n - 1; i++)
    for(int j = 0; j < n - i - 1; j++)
      if(A[j+1] < A[j]) {
        int tmp = A[j + 1];
        A[j + 1] = A[j];
        A[j] = tmp;
      }
}
```

Figure 3: C code of BubbleSort.

to $\beta((c,d)))$, but the path following the `if` branch can be executed at most $\mathbf{min}\{\mathbf{max}\{0,\underline{n}\},3\}$ times. So the asymptotic complexity for the program is $\mathcal{O}(\underline{n})$, but $\mathcal{O}(1)$ for the `if` branch inside the loop.

In the previous example, we have seen how our algorithm works with different branches and more bounds for some edges. Another feature distinguishing our approach from the others is the handling of nested loops. For the BubbleSort of Figure 3, our algorithm computes the bound $\mathbf{max}\{0,\underline{n}-\underline{i}-1\}$ for all edges of the inner loop in the context of the flowgraph induced by the outer loop. After summarization of the outer loop, the bound changes to $\mathbf{max}\{0,\underline{n}-\kappa_1-1\}$, where κ_1 is the path counter of the only backbone of the outer loop. This leads to the `else` branch at line 15 of `processLoop` and the overall bound $\sum_{K=0}^{\mathbf{max}(-1,\underline{n}-2)} \mathbf{max}(0,\underline{n}-K-1)$. If $\underline{n} > 0$, it can be further simplified to $\sum_{K=0}^{\underline{n}-2} \underline{n}-K-1 = \frac{\underline{n}\cdot(\underline{n}-1)}{2}$.

5 Experimental Evaluation

We implemented our algorithm in an experimental program analysis tool called LOOPERMAN. It is built on top of the symbolic execution package BUGST [27] and it intensively uses the SMT solver Z3 [32].

We compare LOOPERMAN with four state-of-the-art loop bound analysis tools: LOOPUS [23], KOAT [7], PUBS [1], and RANK [3]. These freely available tools were mutually compared before and results from their elaborate evaluation are also freely available [29, 30]. We run LOOPERMAN on the same benchmarks originally collected from literature. More precisely, we use the benchmarks translated to C programs by the authors of LOOPUS. We ignore programs with recursive function calls as our tool does not support recursion. In order to make manual inspection of benchmarks and results from their analysis manageable, we also ignore all benchmarks associated with the termination proving tool T2 [6]. At the end, we used 199 benchmarks. They are small C programs (ranging from 7 to 451 lines of code, 26 lines in average) containing various kinds of loops.

We focus on two kinds of bounds: asymptotic complexity bounds for whole programs and *exact* (meaning non-asymptotic) upper bounds for individual program loops. The comparison of exact bounds is restricted to LOOPERMAN and LOOPUS as the other tools use input in different format and (as far as we know) they do not provide mapping of their bounds to the original C code.

We took the results of KoAT, PUBS, and Rank from the mentioned experimental evaluation [29]. In order to obtain the exact bounds for individual loops (which are not publicly available), we run Loopus using a VirtualBox image available on the tool homepage [31]. All experiments use 60 seconds timeout. While KoAT, PUBS, and Rank were executed on a computer with 6GB of RAM and Intel i7 CPU clocked at 3.07 GHz, Loopus and Looperman run on a machine with 8GB of RAM and Intel i5 CPU clocked at 2.5GHz. We believe that this difference is not substantial as all the tools usually either answer very quickly or they fail to determine any bound. The Looperman tool (both sources and Windows binaries), the 199 benchmarks, and all measured data are available at [28].

5.1 Comparison of asymptotic bounds

First we focus on asymptotic complexity bounds for whole programs, which can be produced by all considered tools. Looperman derives them from the upper bounds on edges: each upper bound is transformed into an asymptotic one (e.g. $\mathbf{max}(0, \underline{a}^2 + \underline{b} - 1)$ is transformed into $\mathcal{O}(n^2)$) and maximal asymptotic bound over edges is then taken.

The results are represented in Table 1(a). For each tool, the count of all 199 benchmarks is split into the first three columns. The columns 'correct' and 'incorrect' represent counts of inferred correct and incorrect asymptotic bounds, respectively. The column 'fail' shows counts of benchmarks for which the computation of an asymptotic bound failed. The column 'TO' says how often the fail is due to timeout.

We see that Looperman placed the third after Loopus and KoAT in the count of successfully computed correct asymptotic bounds. If we consider incorrect bounds, Looperman placed the first together with Loopus. Note that the incorrect bounds are due to errors in implementation, the algorithms followed by the tools are in principle sound.

Table 1(b) compares results of Looperman (marked as L.) relatively to each other tool. We can see the number of cases where the other tool succeeds to infer a correct bound while Looperman fails, or vice versa. Further, we can see the number of cases where both tools infer a correct bound and the bound of the other tool is tighter, less tight, or the same as the Looperman's bound.

An importation observation from this table is that for each tool there are benchmarks where the tool fails to compute a correct bound while Looperman succeeds.

5.2 Comparison of exact bounds per loop

Now we compare accuracy of exact upper bounds for individual loops in benchmarks. Note that Looperman computes a loop bound as the sum of bounds on the edges leading from a loop entry node into the loop.

The benchmarks contain 313 loops. Table 2(a) provides for both tools the numbers of inferred correct and incorrect loop bounds, and the number of loops for which no bound is inferred. Table 2(b) compares quality of the inferred loop bounds. It presents the number of loops where one tool produces a correct loop bounds while the other does not, the number of loops where one tool provided asymptotically tighter loop bound than the other, and the number of loops

	correct	incorrect	fail	TO
LOOPERMAN	129	0	70	5
LOOPUS	162	0	37	0
KOAT	140	2	57	22
PUBS	85	29	85	1
RANK	26	2	171	0

(a)

	LOOPUS	KOAT	PUBS	RANK
succeeds, L. fails	38	36	13	7
fails, L. succeeds	5	25	57	110
tighter than L.	10	4	2	0
less tight than L.	8	9	9	1
same bounds	107	91	61	18

(b)

Table 1: Numbers of correct/incorrect/unknown asymptotic bounds for each tool (a) and comparison of the bounds detected by LOOPERMAN (L.) and by other tools (b).

where one tool produces a tighter bound than the other tool, but the difference is not asymptotic (e.g. n versus $2n$). To complete the presented data, let us note that both tools inferred exactly the same bound for 143 loops.

The data in the Table 2(a) show the primary weakness of LOOPERMAN: It gives up too often.[2] On the other hand, data in the Table 2(b) show that LOOPERMAN is able to compute tighter bounds than LOOPUS for interesting number of benchmarks.

The results show that LOOPUS can infer bounds for slightly more loops than LOOPERMAN. However, there are also loops bounded by LOOPERMAN and not by LOOPUS. The biggest advantage of LOOPERMAN is definitely the tightness of its bounds: LOOPERMAN found a tighter bound for 28% of 216 loops bounded by both tools, while LOOPUS found a tighter bound only for 6% of these loops.

6 Related Work

There are several popular approaches to the computation of upper bounds on the number of loop iterations. Notable are especially those based on construction and solving recurrence equations [1, 2, 20, 5], loop iteration counters [5, 15], and ranking functions [3, 7, 24, 23, 26]. Some tools utilize more than one analysis approach.

Recurrence equations play an important role in PUBS [1]. It uses input generated from Java bytecode and computes upper bounds as a non-recursive representation of a given set of recurrent equations obtained from other analysers, like [2]. Another representative work is R-TUBOUND [20]. It rewrites

[2] LOOPERMAN calls the SMT solver Z3 to decide satisfiability of formulae from an undecidable logic. However, none of the fails are due to Z3 as it is able to solve all queries coming from the benchmarks.

	correct	incorrect	fail
LOOPERMAN	227	0	86
LOOPUS	267	3	43

(a)

	LOOPERMAN	LOOPUS
succeeds, the other fails	11	51
asymptotically tighter	16	11
tighter (not asymptotically)	44	2

(b)

Table 2: Numbers of correct/incorrect/unknown loop bounds for LOOPERMAN and LOOPUS (a) and comparison loop bound quality (b).

multi-path loops into single-path ones using SMT reasoning. Obtained loops are then translated into a set of recurrence equations over program variables. Upper bounds are solutions of the equations. This is done by a pattern-based recurrence solving algorithm. The analysis implemented in ABC [5] combines the use of recurrence equations and iteration counters. Authors focus on nested loops there. An inner loop is instrumented by an artificial variable increased by one in each iteration (the counter). Recurrence equations over this variable and regular variables of the loop are constructed and solved (to get their non-recurrent versions). Bounds on the artificial variable are obtained by replacing regular variables by their greatest values in the computed equations. An advanced use of counters can be found in SPEED [15]. Counters are artificial variables instrumented into the analysed program initialised to 0 and incremented by one on back-edges of program loops. A linear invariant generation tool is then used to compute bounds on the counter variables at appropriate locations.

In our approach we also use recurrent equations. They are constructed during symbolic execution and solved in the procedure `computeSummary` as functions of path counters. Nevertheless, our computation is simpler than in the other approaches. And the purpose of equations and their solutions is also different. We use solutions in the process of construction of necessary conditions, whose combinations are latter used for inference of upper bounds.

Similarly to [5] and [15], we also use counters. But in contrast to [5], we introduce a path counter for each acyclic path in the loop instead of one counter for the whole loop. Our approach is thus more related to [15], where counters can also be associated to individual paths (namely, to back-edges of those paths). Nevertheless, we do not perform instrumentation of the counter. Another important difference is that we do not compute bounds on path counters explicitly, while [15] computes them using an invariant generation tool.

A ranking function was originally used for the termination proving [9, 6, 22, 13, 12]. The ranking function is usually defined by assigning an expression over program variables to each program location such that the expression values are always non-negative and decreasing between every two subsequent visits of the particular location. It is apparent that a ranking function can be also used for a computation of upper bounds on the number of iterations of program

loops. RANK [3] reuses results from the termination analysis of a given program, namely an inferred (global) ranking function, to get an asymptotic upper bound on length of all possible executions of that program. KOAT [7] employs a similar approach, but it uses ranking functions of already processed loops to compute bounds on values of program variables. Absolute values of these bounds may in turn improve ranking functions of subsequent loops. LOOPUS [26] transforms an analysed program in particular locations such that program variables appearing there represent ranking functions. There are several heuristics used for this purpose. Program loops are then summarised which is done per individual loop paths. Moreover, a contextualization technique is used to infer which loop paths can be executed in a sequence. The approach was further improved in several directions. Method in [23] merges nested loops in order to compute bounds for programs, where an inner loop affects termination of the outer loop. A technique for computing sizes of variables after loops is introduced in [24].

Our algorithm does not use ranking functions. However, the passing of information from a preceding to a subsequent loop we see in [7] or [24] happens also in our approach. While [7] and [24] implement this via ranking functions, in our algorithm it is available due to the use of the symbolic execution. The loop summarization per individual loop paths presented in [26] is similar to summarization we perform in the procedure `computeSummary`. However, while [26] computes summary as a transitive hull expressed in the domain of a size-change abstraction, we compute precise symbolic values of variables after the loop. On the other hand, we do not consider what loop paths can be executed in a sequence. And finally, in comparison to [23], we do not merge nested loops.

There are other important techniques computing upper bounds, which are, however, less related to our work. For instance, SWEET [17] uses abstract interpretation [10] to derive bounds on values of program variables. The approach further uses a pattern matching for easy and fast processing of loops with a recognized structure. In [16] authors propose an analysis which is able to compute an upper bound on the number of visits of a given program location. The program is transformed such that its execution starts at the location and correctly terminates once the execution reaches the location again. Loops in the transformed program are summarised using an abstract interpretation based iterative algorithm into disjunctive invariants. Upper bounds are computed from the invariants using a non-iterative proof-rules based technique.

Another interesting approach for computation of worst-case complexity, called `WISE` [8], attempts to compute an input for which the execution of a given program is the longest. From symbolic executions of all paths up to a given maximal input size there is inferred a recipe (branch policy generators) which then restricts symbolic execution of all remaining program paths (for unrestricted input size) to the longest ones. The worst-case execution input is then generated from the longest executed path (from its path condition using an SMT solver).

There are further approaches focusing on declarative languages. For example, resource aware ML [18] computes an amortized complexity for recursive functional programs with inductive data types without a consideration of upper bounds dependent on integer values. We can further find numerous techniques for complexity analysis of term rewriting and logic programming [4, 14, 11, 21].

7 Future Work

According to results from the evaluation we will primarily focus on improving scalability of the proposed approach. We have already analysed problematic benchmarks and we have uncovered a couple of promising directions for improvements. Here we present three of them.

1. The loop condition directly implies only inequalities which are non-linear with respect to the path counters. For example, the analysis of the loop `while(x*x<n){ x=x+1; }` leads to the inequality $(\underline{x}^2+\kappa)^2 < \underline{n}$. In order to compute a bound from this constraint we need a more powerful inequation solver.

2. We can add more rules to Algorithm 2 in order to cover more cases of variable value's progression. E.g. in the following loop, the value of `x` would be $\underline{x} + \sum_{i=0}^{\kappa-1}(\underline{y}+1+i)$.

```
while(x<n) {
  y=y+1;
  x=x+y;
}
```

   We plan to extend Algorithm 2 to handle a variety of non-linear progressions.

3. There are further loops where arithmetic progressions of control variables in outer loops are altered in nested loops as illustrated by the following example:

```
while(x>0) {
  x=x-1;
  while(y>0) {
    y=y-1; x=x+1;
  }
}
```

   We plan to deal with such situations according to the solution suggested in [23].

We see no principal issues in the integration of the described improvements into the algorithm presented here. Nevertheless, some changes in the control-flow of the Algorithm 4 and extensions of Algorithms 3 and 2 will be necessary.

Another direction of our future research is an extension to compute also lower bounds. After that, we plan tu use the method for assertion proving and invariant generation.

8 Conclusion

We presented an algorithm computing upper bounds for execution counts of individual instructions of an analysed program during any program run. The

algorithm is based on symbolic execution and the concept of path counters. The upper bounds are parametrized by input values of the analysed program. Evaluation of our experimental tool LOOPERMAN shows that our approach is slightly less robust than the leading loop bound analysis tools LOOPUS and KOAT (i.e. it infers a bound in less cases). On the positive side, the loop bounds detected by LOOPERMAN are often tighter than these found by other tools, which may be a crucial advantage in some applications including the worst case execution time (WCET) analysis.

References

[1] E. Albert, P. Arenas, S. Genaim, and G. Puebla. Automatic inference of upper bounds for recurrence relations in cost analysis. In *Proceedings of the 15th International Symposium on Static Analysis*, Lecture Notes in Computer Science, pages 221–237. Springer-Verlag, 2008.

[2] E. Albert, P. Arenas, S. Genaim, G. Puebla, and D. Zanardini. Cost analysis of java bytecode. In *Programming Languages and Systems*, volume 4421 of *Lecture Notes in Computer Science*, pages 157–172. Springer Berlin Heidelberg, 2007.

[3] C. Alias, A. Darte, P. Feautrier, and L. Gonnord. Multi-dimensional rankings, program termination, and complexity bounds of flowchart programs. In *Proceedings of the 17th International Conference on Static Analysis*, Lecture Notes in Computer Science, pages 117–133. Springer-Verlag, 2010.

[4] M. Avanzini and G. Moser. A combination framework for complexity. In *Proceedings of 24th International Conference on Rewriting Techniques and Applications*, volume 21 of *Leibniz International Proceedings in Informatics*, pages 55–70, 2013.

[5] R. Blanc, T. A. Henzinger, T. Hottelier, and L. Kovács. ABC: Algebraic bound computation for loops. In E. M. Clarke and A. Voronkov, editors, *Proceedings of the 16th International Conference on Logic for Programming, Artificial Intelligence, and Reasoning*, volume 6355 of *Lecture Notes in Computer Science*, pages 103–118. Springer-Verlag, 2010.

[6] M. Brockschmidt, B. Cook, and C. Fuhs. Better termination proving through cooperation. In *Computer Aided Verification*, volume 8044 of *Lecture Notes in Computer Science*, pages 413–429. Springer Berlin Heidelberg, 2013.

[7] M. Brockschmidt, F. Emmes, S. Falke, C. Fuhs, and J. Giesl. Alternating runtime and size complexity analysis of integer programs. In *Tools and Algorithms for the Construction and Analysis of Systems*, volume 8413 of *Lecture Notes in Computer Science*, pages 140–155. Springer Berlin Heidelberg, 2014.

[8] J. Burnim, S. Juvekar, and K. Sen. WISE: Automated test generation for worst-case complexity. In *Proceedings of the 31st International Conference on Software Engineering*, pages 463–473. IEEE Computer Society, 2009.

[9] B. Cook, A. Podelski, and A. Rybalchenko. TERMINATOR: Beyond safety. In *Proceedings of the 18th International Conference on Computer Aided Verification*, volume 4144 of *Lecture Notes in Computer Science*, pages 415–418. Springer-Verlag, 2006.

[10] P. Cousot and R. Cousot. Abstract interpretation: a unified lattice model for static analysis of programs by construction or approximation of fixpoints. In *Conference Record of the Fourth Annual ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages*, pages 238–252. ACM Press, New York, NY, 1977.

[11] S. K. Debray and N.-W. Lin. Cost analysis of logic programs. *ACM Transactions on Programming Languages and Systems, TOPLAS'93*, 15(5):826–875, Nov. 1993.

[12] S. Falke, D. Kapur, and C. Sinz. Termination analysis of imperative programs using bitvector arithmetic. In *Proceedings of the 4th International Conference on Verified Software: Theories, Tools, Experiments*, Lecture Notes in Computer Science, pages 261–277. Springer-Verlag, 2012.

[13] C. Fuhs, J. Giesl, M. Plücker, P. Schneider-Kamp, and S. Falke. Proving termination of integer term rewriting. In *Rewriting Techniques and Applications*, volume 5595 of *Lecture Notes in Computer Science*, pages 32–47. Springer Berlin Heidelberg, 2009.

[14] J. Giesl, T. Ströder, P. Schneider-Kamp, F. Emmes, and C. Fuhs. Symbolic evaluation graphs and term rewriting: A general methodology for analyzing logic programs. In *Proceedings of the 14th Symposium on Principles and Practice of Declarative Programming*, Lecture Notes in Computer Science, pages 1–12. ACM, 2012.

[15] S. Gulwani, K. K. Mehra, and T. Chilimbi. SPEED: Precise and efficient static estimation of program computational complexity. In *Proceedings of the 36th Annual ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages*, pages 127–139. ACM, 2009.

[16] S. Gulwani and F. Zuleger. The reachability-bound problem. In *Proceedings of the 31st ACM SIGPLAN Conference on Programming Language Design and Implementation*, pages 292–304. ACM, 2010.

[17] J. Gustafsson, A. Ermedahl, C. Sandberg, and B. Lisper. Automatic derivation of loop bounds and infeasible paths for WCET analysis using abstract execution. In *Real-Time Systems Symposium, 2006. RTSS '06. 27th IEEE International*, pages 57–66, 2006.

[18] J. Hoffmann, K. Aehlig, and M. Hofmann. Multivariate amortized resource analysis. *ACM Transactions on Programming Languages and Systems*, 34(3):14:1–14:62, 2012.

[19] J. C. King. Symbolic execution and program testing. *Communications of the ACM*, 19(7):385–394, 1976.

[20] J. Knoop, L. Kovács, and J. Zwirchmayr. Symbolic loop bound computation for WCET analysis. In *Proceedings of the 8th International Conference on Perspectives of System Informatics*, Lecture Notes in Computer Science, pages 227–242. Springer-Verlag, 2012.

[21] J. Navas, E. Mera, P. López-García, and M. Hermenegildo. User-definable resource bounds analysis for logic programs. In *Logic Programming*, volume 4670 of *Lecture Notes in Computer Science*, pages 348–363. Springer Berlin Heidelberg, 2007.

[22] A. Podelski and A. Rybalchenko. ARMC: The logical choice for software model checking with abstraction refinement. In *Proceedings of the 9th International Conference on Practical Aspects of Declarative Languages*, Lecture Notes in Computer Science, pages 245–259. Springer-Verlag, 2007.

[23] M. Sinn, F. Zuleger, and H. Veith. A simple and scalable static analysis for bound analysis and amortized complexity analysis. In *Proceedings of the 16th International Conference on Computer Aided Verification - Volume 8559*, Lecture Notes in Computer Science, pages 745–761. Springer-Verlag New York, Inc., 2014.

[24] M. Sinn, F. Zuleger, and H. Veith. Difference constraints: An adequate abstraction for complexity analysis of imperative programs. In *Proceedings of the Formal Methods in Computer-Aided Design*, pages 144–151, 2015.

[25] J. Strejček and M. Trtík. Abstracting path conditions. In *Proceedings of the 2012 International Symposium on Software Testing and Analysis*, pages 155–165. ACM, 2012.

[26] F. Zuleger, S. Gulwani, M. Sinn, and H. Veith. Bound analysis of imperative programs with the size-change abstraction. In *Proceedings of the 18th international conference on Static analysis*, pages 280–297. Springer-Verlag, 2011.

[27] `Bugst`. `http://sourceforge.net/projects/bugst/`.

[28] `Looperman, benchmarks, and evaluation`. `https://sourceforge.net/projects/bugst/files/Looperman/1.0.0/`.

[29] `Evaluation results of KoAT, PUBS, and RANK`. `http://aprove.informatik.rwth-aachen.de/eval/IntegerComplexity`.

[30] `Loopus: Comparison to the tools KoAT, PUBS and Rank`. `http://forsyte.at/static/people/sinn/loopus/CAV14/index.html`.

[31] `Loopus`. `http://forsyte.at/software/loopus/` version from February 2, 2015.

[32] `Z3`. `https://github.com/Z3Prover/z3`.